TRANSMITTAL OF APPEAL BRIEF (Large Entity)

Docket No. 9261(NCR.0021US)

In Re Application Of: **M. Cameron Watson et al.**

Serial No.	Filing Date	Examiner	Group Art Unit
09/728,852	12/01/2000	H. Mahmoudi	2175

Invention: **METHOD AND APPARATUS FOR PATITIONING DATA FOR STORAGE IN A DATABASE**

TO THE COMMISSIONER FOR PATENTS:

Transmitted herewith in triplicate is the Appeal Brief in this application, with respect to the Notice of Appeal filed on

The fee for filing this Appeal Brief is: **$320.00**

☐ A check in the amount of the fee is enclosed.

☐ The Director has already been authorized to charge fees in this application to a Deposit Account.

☒ The Director is hereby authorized to charge any fees which may be required, or credit any overpayment to Deposit Account No. **50-1673 (9261)**

Dated: 9-16-03

Signature

Dan C. Hu
Reg. No. 40,025
TROP, PRUNER & HU, P.C.
8554 Katy Freeway
Suite 100
Houston, Texas 77024

I certify that this document and fee is being deposited on 9-16-2003 with the U.S. Postal Service as first class mail under 37 C.F.R. 1.8 and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Dawn L. Thomas

Signature of Person Mailing Correspondence

Dawn L. Thomas

Typed or Printed Name of Person Mailing Correspondence

cc:

#10
10/2/03
A.W.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Applicant: M. Cameron Watson et al.	§	Examiner: H. Mahmoudi
Serial No.: 09/728,852	§	Group Art Unit No: 2175
Filed: December 1, 2000	§	Atty Dkt. No. NCR.0021US (9261)
Title: METHOD AND APPARATUS FOR PATITIONING DATA FOR STORAGE IN A DATABASE	§	

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

Sir:

Applicant respectfully appeals from the final rejection mailed April 4, 2003.

I. REAL PARTY IN INTEREST

The real party in interest is NCR Corporation, the assignee of the present application by virtue of the assignment recorded at Reel/Frame 011344/0971.

II. RELATED APPEALS AND INTERFERENCES

None.

III. STATUS OF THE CLAIMS

Claims 1-9 and 11-31 have been finally rejected and are the subject of this appeal.

IV. STATUS OF AMENDMENTS

No amendments were submitted after the final rejections.

Date of Deposit: September 16, 2003
I hereby certify under 37 CFR 1.8(a) that this correspondence is being deposited with the United States Postal Service as **first class mail** with sufficient postage on the date indicated above and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313.

Dawn L. Thomas

V. SUMMARY OF THE INVENTION

Generally, when data is transmitted from one system to another, such as from a server system to a client system, a certain amount of efficiency in retrieving the data is desired. Often, a database may be accessed by a plurality of users or one user may issue a plurality of queries. Allowing for more efficient access of data from a database can ease the computing congestion that occurs when multiple queries are being processed. Specification p. 3, ll. 22-26.

One way to improve performance in extracting data from a database is to set up a database such that data is distributed across a plurality of servers, called data servers (e.g. a distributive or parallel database system). Partitioning data stored in the database in an efficient manner will improve the utilization of data servers in accessing data in the database. Some embodiments of the invention provide a method and apparatus for partitioning related data so that a reduced number of data servers are invoked in response to a query to access the data. Specification p. 3, l. 27-p. 4, l. 8. One way of achieving the goal of invoking a smaller number of data servers is by partitioning data in a database based on certain user inputs. Specification, p. 4, ll. 20-31.

In one implementation, depicted in Fig. 2 of the specification, a user may direct a client system 130 to send information relating to characteristics associated with data to be stored in a database system 120 to a database controller 110. The database controller 110 utilizes this information to partition the data for efficient storage and access. Specification p. 5, ll. 13-19. The database system 120 includes data servers 260 (Fig. 2) that are capable of accessing corresponding storage locations in the database system 120 to store and retrieve data. Specification p. 5, ll. 20-22. In response to data queries, the database controller 110 chooses one or plural data servers 260 within the database system 120 to search for the data that is being requested. Specification p. 5, l. 23-p. 6, l. 3.

Using a partitioning algorithm, a partitioner 340 (Fig. 3) in the database controller 110 is capable of grouping related data onto one or a limited number of data servers 260 (Fig 2). This can reduce the need to search a large number of data servers 260 to find requested data. The partitioner 340 then stores the related data portions in storage units that are under the supervision of the one or a limited number of data servers 260. Information utilized by the partitioner 340 for performing grouping of data sets is stored in a partitioner data storage 350 (Fig 3). Specification p. 6, ll. 12-20.

In one implementation, the partitioning algorithm is based upon a straight-line segment approximation of partitioning values of a given data set. The partitioner 340 uses the partitioning values to determine which data servers 260 are to be invoked to perform a data access operation. Specification p. 6, l. 21-p. 7, l. 2.

A simplified illustrative example of a data search in accordance with the present invention is described below. For example, using the partitioning method described above, data relating six colors, such as red, orange, yellow, green, blue, and purple, are stored by six respective data servers 260. In response to a data query that requests information for the color reddish-orange, the database controller 110 performs a data search narrowing function using the organizational information stored in the partitioner data storage 350. The partitioner 340 is then able to make a determination that two of the data servers 260 contain data related to the colors red and orange. Therefore, the query coordinator 320 selects the two data servers 260 (instead of all six data servers) that can locate the requested information. Specification p. 8, l. 26-p. 9, l. 8.

In one arrangement, if the data search that is performed in response to the data query does not result in the finding of the data corresponding to reddish-orange color, the field of search is broadened. The field of search is broadened to include other data servers 260 that may have

addresses that will lead to data corresponding to the color reddish-orange in the database system 120. With the acquisition of data correlating to the colors red and orange, a determination is then made regarding whether the desired data has been acquired. Once the data search results are determined to be accurate in relation to the query search within a predetermined confidence level, the search by the database controller 110 is terminated. The acquired data is sent back to the requester (e.g., client system 130). Specification p. 9, ll. 9-21.

Although several embodiments have been described above, other embodiments are also covered by the claims on appeal.

VI. ISSUES

A. Can References That Do Not Teach or Suggest The Subject Matter Of Independent Claims 1, 11, and 19 Render Obvious Claims 1-3, 6-9, 11-23, 25-27, 29, and 30?

B. Can References That Do Not Teach or Suggest The Subject Matter of Claims 24, 28, and 31 Render Those Claims Obvious?

C. Can References That Do Not Teach or Suggest The Subject Matter Of Claims 4 and 5 Render Those Claims Obvious?

VII. GROUPING OF THE CLAIMS

Group 1: Claims 1-3, 6-9, 11-23, 29, and 30 can be grouped together.

Group 2: Claims 24, 28, and 31 can be grouped together.

Group 3: Claims 4 and 5 can be grouped together.

Within each group, the claims stand and fall together.

VIII. ARGUMENT

All claims should be allowed over the cited references for the reasons set forth below.

A. Can References That Do Not Teach or Suggest The Subject Matter Of Independent Claims 1, 11, and 19 Render Obvious Claims 1-3, 6-9, 11-23, 25-27, 29, and 30?

Independent claim 1 was rejected as being obvious over the hypothetical combination of Reiner (U.S. Patent No. 6,289,334) and Nori (U.S. Patent No. 6,061,690). Even if Reiner and Nori can be properly combined, they do not teach or suggest the recited combination of elements in claim 1. The Examiner incorrectly cited to columns 2, 3, and 25 of Reiner as teaching the last clause of claim 1, namely "in response to a database query, selecting less than all the plural data servers based on the partitioning of the data to reduce a number of data servers involved in processing the database query." 4/4/03 Office Action at 3. The passage in column 2, lines 63-65, refers to a "standard" interface that is accessed by clients which are the sources of queries. The passage in column 3, at lines 32-52, refers to an aspect of the purported invention of Reiner that relates to the use of a "driving table" whose partitions are targeted by subqueries (decomposed from a main query). The passage at column 25, lines 39-48, refers to the problem of partition skew that results in unequal-sized partitions. The column 25 passage states that during a latter part of a query execution, and possibly even during the entire query, some partitions will have no more rows to fetch, which reduces the degree of parallelism for the remainder of the query.

None of these passages teach or suggest *selecting* less than all the plural data servers (in response to the database query) based on the partitioning of the data to reduce a number of data servers involved in processing the database query. The passage in claim 25 cited by the Examiner states that during a latter part of a query execution (and sometimes during the entire query), some partitions will have no rows to fetch. To address this problem, Reiner proposes

that the database query be *intercepted*, with the intercepted query *decomposed* into multiple sub-queries. Reiner, 2:65-3:3. The sub-queries are submitted in parallel to access respective partitions. Reiner, 7:64-8:40. An example of the sub-queries is shown in lines 25-30 of column 9, in which partition numbers are actually included in each sub-query for indicating the partition that the sub-query is to operate on. In the column 9 example of Reiner, Applicant notes that *all* partitions are selected in response to the intercepted query, since multiple sub-queries are decomposed from the intercepted query to access the partitions in parallel.

The goal of Reiner is to enhance parallelism in executing *one* query (the intercepted query) by decomposing the intercepted query into multiple sub-queries that operate on respective partitions. *See* Reiner, 2:31-36 ("Unfortunately, prior art DBMS's have not proven capable of taking full advantage of the power of such multiprocessing systems and, particularly, their power to simultaneously process data (in parallel) from multiple partitions on multiple storage devices with multiple central processing units."). As further stated by Reiner, "the bucket size is chosen to insure that hash buckets are spread over storage devices to maximize the potential for parallel access." Reiner, 10:48-50. Moreover, Reiner states that "the database 72 is organized to achieve the best mix of I/O parallelism and hit ratio" since greater I/O parallelism means that more threads can be used, in parallel, to initiate data retrievals. Reiner, 10:62-67. In view of the foregoing, it is respectfully submitted that Reiner does not teach or suggest selecting (in response to a database query) less than all the plural data servers as recited in claim 1.

In fact, Reiner teaches *the opposite*. To achieve the goal of taking full advantage of the power of multiprocessing systems, all processing elements have to be selected in the Reiner system by issuing multiple sub-queries (decomposed from an intercepted query) to the database

system. Therefore, a person of ordinary skill in the art looking to the teachings of Reiner would have been led away from the claimed invention, rather than towards it.

Nori also fails to teach or suggest this missing element of claim 1. Therefore, even if the teachings of Reiner and Nori can be combined, the combination does not teach or suggest the claimed invention. Therefore, for at least this reason, a *prima facie* obviousness rejected has not been established with respect to claim 1.

In addition, Applicant notes another error in the obviousness rejection of claim 1. The Examiner stated that "a set of input parameters" mentioned in column 11, lines 25-37, of Nori, constitutes the element "information associated with at least one characteristic of the data" recited in claim 1. 4/4/03 Office Action at 3. The cited passage of Nori refers to functions each receiving a set of parameters, with the set of function names and their parameters defining an application programming interface (API) to the functions. Nori, 11:33-35. There is no indication that the input parameters to the functions constitute "information associated with at least one characteristic of the data." Furthermore, and more significantly, the "receiving" act of claim 1 cannot be read in isolation—the claim further recites that the information associated with the at least one characteristic of the data is used to *partition* data for storage in storage units associated with plural data servers. There is no absolutely no suggestion whatsoever in Nori that the set of parameters is used for partitioning data. Thus, contrary to the assertion made by the Examiner, Nori fails to teach or suggest "receiving information associated with at least one characteristic of the data," where partitioning of the data is "based on the characteristics associated with the data." For this further reason, the Examiner has failed to establish *prima facie* obviousness rejection with respect to claim 1 over the hypothetical combination of Reiner and Nori.

The Examiner's obviousness rejection is further defective because there is no suggestion or motivation to combine Reiner and Nori. Reiner actually teaches away from the invention. Reiner notes limitation associated with prior system in which processing of partitions occur even though some partitions have no more rows to fetch. Reiner, 25:39-47. To address this problem, Reiner proposes a mechanism that intercepts a query and decomposes the intercepted query into plural sub-queries that are executed in parallel to enhance parallelism. The goal of Reiner is to enhance parallelism by executing the sub-queries in parallel so that the partitions can be accessed in parallel. This goal of Reiner is contrasted with claim 1, which recites "in response to a database query, selecting *less than all* the plural data servers based on partitioning of the data to *reduce a number of data servers* involved in processing the database query." Because Reiner teaches away from the invention, there can be no motivation or suggestion to combine Reiner and Nori in the manner proposed by the Examiner.

The Examiner has also failed to establish how the disparate teachings of Reiner and Nori can be combined. Reiner focuses on intercepting a main query and decomposing the main query into sub-queries for submission to a database management system. Reiner, 2:61-3:3. Nori refers to functions of an API, and the set of parameters associated with each function. The teachings of Reiner and Nori are unrelated to each other. It is unclear to Applicant, and the Examiner has provided no rationale to justify why the teaching in Reiner with respect to decomposing a main query into plural subqueries can be combined with the teaching in Nori regarding the functions and parameters of an API. The Examiner stated that "[i]t would have been obvious to a person having ordinary skill in the art at the time the invention was made to have modified Reiner et al by the teaching of Nori et al, because receiving information associated with at least one characteristic of the data, would enable the user to receive relational information about the data

and define the relationships between data segments partitioned and stored in the plurality of storage devices." 4/4/03 Office Action at 3-4. The cited motivation to combine the references is clearly defective, because the input parameters to functions of an API disclosed by Nori have nothing to do with the decomposition of a main query into sub-queries disclosed in Reiner. The Examiner has failed to explain how the API function input parameters can be used in the Reiner system. Applicant respectfully submits that there simply is no reason to incorporate the API function input parameters of Nori into the Reiner query decomposition mechanism. In fact, the Examiner has failed to establish if the API function input parameters of Nori can even technically be incorporated into the query decomposition mechanism.

The obviousness rejection of claim 1 is defective for this further reason. Independent claims 11 and 19 are allowable for reasons similar to those given above with respect to claim 1.

For the reasons above, the final rejection of claims 1-3, 6-9, 11-23, 25-27, 29, and 30 should be reversed.

B. Can References That Do Not Teach or Suggest The Subject Matter of Claims 24, 28, and 31 Render Obvious Thos Claims?

Claims 24, 28, and 29 depend from independent claims 1, 11, and 19, respectively, and thus are allowable for at least the same reasons as corresponding independent claims.

Moreover, these claims are further allowable for the following reasons. Claim 24 recites that the method of claim 1 further comprises determining whether search results are satisfactory, and selecting at least one more data server to process the database query if the search results are not satisfactory. The Examiner cited two passages in columns 16 and 18 of Reiner as teaching the determining act, and the Examiner cited to a passage in column 19 as teaching the selecting act. 4/4/03 Office Action at 11. The cited passages do not teach or suggest the recited elements

of claim 24. The passages in columns 16 and 18 of Reiner describe the need for a combining query or function to produce the correct results for the original received query. The column 19 passage of Reiner states that queries containing the aggregate functions STDDEV or VARIANCE can be effectively parallelized through target list modification and a sophisticated combining query. Note that claim 24 recites that the selection of at least one more data server to process the database query is performed if the search results are determined not to be satisfactory. This causal relation between the determining and selecting acts of claim 24 is not described or suggested at all in Reiner. Claim 24 is allowable over the hypothetical combination of Reiner and Nori for this further reason.

The other claims, 28 and 31, are similarly allowable.

For the reasons above, the final rejection of claims 24, 28, and 31 should be reversed.

C. Can References That Do Not Teach or Suggest The Subject Matter Of Claims 4 and 5 Render Thos Claims Obvious?

Claim 4 depends from claim 3, which depends from claim 1. Therefore, claim 4 is allowable for at least the reasons as claim 1.

Claim 4 is further allowable over the hypothetical combination of Reiner, Nori, and Natarajan for the following reasons. As conceded by the Examiner, the combination of Reiner and Nori does not teach "defining straight-line segments based on at least one of the average value of the data, the uniform distribution of the data, the minimum value of the data, and the maximum value of the data." 4/4/03 Office Action at 14. Rather, the Examiner cited to Natarajan as teaching the elements recited by claim 4. The Examiner's reliance on Natarajan as teaching this element of claim 4 is mis-placed. Natarajan describes a filtering system and method for reducing random noise in signals in modern digital and analog systems. Natarajan,

1:8-13. "According to Natarajan, its "invention provides a general method and a system for reducing random noise in signal using data compression according to the principle of Occam's Razor." Natarajan, 2:27-30. Natarajan is completely unrelated to the recited elements of the claim 4, where partitioning of data includes defining straight-line segments based on one of several factors. There simply is no motivation or suggestion to combine the references in the manner asserted by the Examiner.

The Examiner fails to discuss how the teachings of Natarajan can be applied to the teachings of Reiner and/or Nori. Furthermore, even if the three references can be combined, such combination does not teach or suggest the element of claim 4, namely, that partitioning of data comprises defining straight-line segments based on at least one of several factors.

For the reasons set above, the final rejection of claims 4 and 5 should be reversed.

IX. CONCLUSION

Applicant respectfully requests that each of the final rejections be reversed and that the claims subject to this appeal be allowed to issue.

Respectfully submitted,

Date: 9-16-03

Dan C. Hu, Reg. No. 40,025
TROP, PRUNER & HU, P.C.
8554 Katy Fwy, Ste 100
Houston, TX 77024-1805
713/468-8880 [Phone]
713/468-8883 [Facsimile]

CLAIMS ON APPEAL

1. A method, comprising:

receiving data to be stored in a database system having plural data servers;

receiving information associated with at least one characteristic of the data;

partitioning the data for storage in the database system based on the characteristic associated with the data; and

storing the partitioned data in storage units associated with the plural data servers; and

in response to a database query, selecting less than all the plural data servers based on the partitioning of the data to reduce a number of data servers involved in processing the database query.

2. The method of claim 1, wherein receiving the information comprises receiving the information from a client system.

3. The method of claim 1, wherein receiving the information comprises receiving at least one of an average value of the data, a uniform distribution of the data, a minimum value of the data, and a maximum value of the data.

4. The method of claim 3, wherein partitioning the data comprises defining straight-line segments based on at least one of the average value of the data, the uniform distribution of the data, the minimum value of the data, and the maximum value of the data.

5. The method of claim 4, wherein partitioning the data further comprises defining breakpoints to provide the straight-line segments.

6. The method of claim 1, wherein partitioning the data for storage in the database system comprises dividing the data into buckets containing related data.

7. The method of claim 1, wherein partitioning the data comprises organizing the data into related portions.

8. The method of claim 7, wherein partitioning the data further comprises executing an algorithm to organize the data.

9. The method of claim 1, wherein storing the partitioned data in the database system comprises storing the partitioned data in a relational database system.

11. A system, comprising:

a database;

a network interface;

plural storage modules and data servers;

a database controller coupled to the database, wherein the database controller is adapted to receive partitioning information and perform a partitioning task on data received through the network interface based on the partitioning information to partition the data into plural groups,

the database controller adapted to further store the plural groups of the data partitioned by the partitioning task into plural storage modules associated with corresponding plural data servers,

the database controller adapted to select, in response to a database query, less than all the plural data servers based on the partitioning information to reduce a number of data servers involved in processing the database query.

12. The system of claim 11, wherein the database is part of a parallel database system.

13. The system of claim 11, wherein the database is a relational database.

14. The system of claim 11, wherein the database controller comprises:

a query coordinator coupled to the network interface, the query coordinator to receive the database query from the network interface;

a partitioner to partition data and perform selecting of less than all the plural data servers; and

a partitioner data storage coupled to the partitioner, the partitioner data storage to store the partitioning information associated with at least one characteristic of the data to enable the partitioner to partition data.

15. The system of claim 14, wherein the partitioner is capable of executing an algorithm, based on the stored partitioning information, for partitioning the data.

16. The system of claim 15, wherein the plural data servers are adapted to store and access partitioned data in the database.

17. The system of claim 11, further comprising a client system, wherein the client system sends data to the database through the network interface.

18. The system of claim 17, wherein the client system is adapted to further send the partitioning information to be used by the database controller to partition the data.

19. An article comprising one or more storage media containing instructions that when executed cause a device to:

receive information associated with at least one characteristic of data to be stored into a database system from a remote device;

partition the data for storage in the database system based on the characteristic of the data;

store the partitioned data in the database system in plural storage modules associated with plural data servers; and

in response to a database query, select less than all the data servers based on the information to reduce a number of data servers involved in processing the database query.

20. The article of claim 19, wherein the instructions when executed cause the device to execute an algorithm to partition the data.

21. The article of claim 19, wherein the instructions when executed cause the device to divide the data into segments containing related data.

22. The method of claim 1, wherein receiving the information comprises receiving organizational information, and wherein selecting less than all the plural data servers is based on the organizational information.

23. The method of claim 22, wherein selecting less than all the plural data servers is based on the organizational information and a characteristic of data requested by the database query.

24. The method of claim 1, further comprising:

retrieving search results obtained by the selected data servers;

determining whether the search results are satisfactory; and

selecting at least one more data server to process the database query if the search results are not satisfactory.

25. The method of claim 1, wherein partitioning the data comprises partitioning the data into logical groups.

26. The method of claim 1, further comprising storing the information by a partitioner, wherein selecting less than all the data select is performed at least in part by the partitioner.

27. The system of claim 11, the database controller to select less than all the plural data servers based on the partitioning information and a characteristic of data requested by the database query.

28. The system of claim 11, wherein the selected data servers are adapted to retrieve search results in response to the database query, and the database controller is adapted to determine whether the search results are satisfactory and to select at least one more data server to process the database query if the search results are not satisfactory.

29. The article of claim 19, wherein the instructions when executed cause the device to receive information comprising partitioning information.

30. The article of claim 29, wherein the instructions when executed cause the device to select less than all the plural data servers based on the partitioning information and a characteristic of data requested by the database query.

31. The article of claim 19, wherein the instructions when executed cause the device to:

retrieve search results obtained by the selected data servers;

determine whether the search results are satisfactory; and

select at least one more data server to process the database query if the search results are not satisfactory.